

DAS

NETZ

digitalization and society
English edition



All hands on deck!

Wow, we're digital! Today, nothing is done without at least some element of digitalization. Smart procedures and digital management are everywhere. By now, everybody knows that the digital transition will play a role in their life, their company, their white paper... This is a good start. And yet, buzzwords alone do not add up to an intelligent strategy. Neither does a revamped approach to automation presented as digitization. Of course, the ubiquitous efforts to shape the digital future are very welcome. But still: these efforts could go even further and be of more consequence.

A digital transition is happening in China. A new tool here, a new idea there—implemented about ten times faster, and in a way that is ten times more encompassing than in Europe. How does one govern a society? How does one manage an economy? How are we to think in comprehensive digital ecosystems instead of limiting the focus to individual products and services? If you need inspiration, one should first look to the East. In this edition, you will get a fascinating view of the Chinese way into the digital age.

The internet and global digitization provide easy and rapid means to look beyond one's horizon—and that is precisely what governments, companies and individuals should be doing. The digital world allows us to exchange perspectives, ideas and concepts, and to learn from each other on a truly global scale. This ease of communication is one of the simplest and most basic aspects of the connected world, and is still one of its major benefits.

In a rapidly changing world, governments should be coordinating and constantly updating their digital agendas. They need to take all of the digital transition's enormous potentials into account—not only regarding the economy and innovation, but also with respect to non-economic aspects. A connected society is simultaneously an ideal sphere within which to discuss the values that shape our future, and within which to establish a balance between the interests of different stakeholders. But it is not only the state that is responsible for this discourse. All citizens are called

on to express themselves. Where the definition and elaboration of values in the digital sector are concerned, almost all of us are sitting in front of the infamous blank page. It is one of our generation's tasks to define what we want and what we don't want, what is desirable and what we as a society deem unacceptable. This includes the continued development of our conventions. We will have to part with some old and familiar principles; we will have to establish new parameters. There is nothing more exciting! So please feel welcome to become political and to participate in this journey. No matter where you are—what counts is that you use your voice.

On behalf of the editorial team
Philipp Otto



Photo: Bettina Volke

Philipp Otto is the founder of the think tank iRights.Lab and the publishing house iRights.Media. He is a publisher of iRights.info. He develops strategies and concepts to successfully shape the digital transition. In doing so, he works both with and for governments, parliaments, companies and representatives of civil society.

Life

- 8** **The internet works in mysterious ways**
- 12** **Predictive healthcare: Medicine in the data revolution**
by Lydia Heller
- 16** **Attack of the fridges** by Jessica Binsch
- 18** **Artificial intelligence: The dreaming algorithm**
by Christoph Drosser
- 24** **What happened online? January 2016**
- 26** **The digital doping hunt** by Martin Einsiedler
- 30** **#scanallfishes**
- 34** **Universities face digital challenges**
by Ada Pellert
- 37** **Digitalization is happening... in your aerobics class**
Interview with Stefan Will
- 40** **What happened online? February 2016**
- 43** **Learning to program is a skill for life**
by Gerhard Seiler and Jutta Schneider
- 46** **Emancipation through citizen science**
by Henry Steinhau
- 50** **If you were Queen of the Internet, what would be your first decree?**
- 53** **Catastrophe! Communication in states of emergency**
by Julia Schönborn
- 56** **What happened online? March 2016**
- 58** **The Twitter troll's digital alter ego**
by René Walter
- 62** **The summer of Pikachu** by Dennis Kogel
- 65** **Barfing unicorns and puppy faces:**
What is the secret ingredient in Snapchat's success?
by Duygu Gezen
- 68** **Really great sex—just do it right** by Christine Olderdissen
- 71** **Happy coincidences and personalized filter bubbles**
by Christoph Lutz
- 74** **What happened online? April 2016**
- 76** **Who runs the internet?**

Politics

- 80** **The 2016 US election campaign: Digital mud-wrestling** by Lukas Schöne
- 84** **The evolution of the digital election**
by Adrian Rosenthal and Axel Wallrabenstein
- 88** **Why social bots threaten our democracy**
by Martin Fuchs
- 92** **This message will self-destruct in three seconds...**
by Aleksandra Sowa
- 96** **What happened online? May 2016**
- 98** **Politics has to be shaped by people**
Interview with Nadine Schön
- 101** **The SPD is becoming more and more digital**
Interview with Katarina Barley
- 105** **Open? Free? Inclusive? Internet governance at the crossroads**
by Henning Lahmann
- 108** **How are the rules of the internet made?**
Interview with Wolfgang Kleinwächter
- 112** **The digital rich-poor divide**
- 114** **Turkey censors both online and on the streets** by Hauke Gierow
- 118** **What happened online? June 2016**
- 121** **Will digitalization destroy our values?**
by Sabine Leutheusser-Schnarrenberger
- 123** **What have algorithms got to do with human rights?**
Interview with Ben Wagner
- 126** **An internet of self-determination, diversity and participation**
Interview with Heiko Maas
- 131** **Missed opportunities, half-hearted solutions** by Halina Wawzyniak
- 132** **Digital policy decisions: Fail!** by Konstantin von Notz
- 134** **Digitalization is like the industrial revolution**
Interview with Christian Lindner
- 136** **We need more European standards**
Interview with Jan Philipp Albrecht
- 139** **Cautious steps into the minefield** by Joerg Heidrich
- 142** **Personalized pricing needs rules** by Klaus Müller
- 146** **What happened online? July 2016**
- 149** **Digital by default** by Julia Kloiber

Economy

- 154 Towards a giant world computer**
by Stefan Mey
- 158 We're about to experience a real killer app for blockchain**
Interview with Shermin Voshmgir
- 160 Bitcoin: The ascent of a borderless currency**
by Imogen Goodman
- 165 The political promises of Bitcoin**
Interview with Andreas M. Antonopoulos
- 168 What happened online? August 2016**
- 170 Chinese internet firms find success with indecent exposure**
by Finn Mayer-Kuckuk
- 173 Wallet-less payment is an everyday affair in China**
- 175 The next Silicon Valley? It could be here.**
by Tobias Schwarz
- 181 On robots and class struggle: Are we being replaced by machines?**
by Mads Pankow
- 184 What happened online? September 2016**
- 187 We can guarantee the availability of the internet**
Interview with Harald Summa
- 191 Political action shouldn't make things worse**
Interview with Alexander Hüsing

The background of the page is a complex, abstract line drawing in a light orange or terracotta color. It consists of numerous overlapping, irregular, and somewhat chaotic lines that create a sense of movement and depth, resembling a stylized map or a network of paths. The lines vary in thickness and density, with some areas being more heavily drawn than others.

Culture

- 196** TV at the crossroads of internet and humanity
by Andreas Busche
- 200** Hatsune Miku, the world's first cybernetic star
by Finn Mayer Kuckuk
- 202** What Pokémon Go has in common with Locative Art
by Valie Djordjevic
- 206** What happened online? October 2016
- 208** The moment when an unexpected perturbation
changes the system from within
Interview with Tatiana Bazzichelli
- 211** The inhibitions of Richard W. by Christian Rickerts
- 214** Digital passport: Citizen Ex
- 216** The myth of struggling through
Interview with Lisa Basten
- 220** What happened online? November 2016
- 222** Caring for customs and heritage of the internet
by Dirk von Gehlen
- 226** So two computers meet, and one says...
by Alard von Kittlitz and Johannes Gernert
- 230** Gamification: The brain's addiction by Ippolita
- 236** Artists of this edition
- 237** About iRights.Media
- 238** Imprint

The internet works in mysterious ways



Chewbacca Mom

Candace Payne decided to document her irrepressible delight in the impulse purchase of a Chewbacca mask, depicting the famous character from the Star Wars films, and share the video on the internet. A few days later, the mask was everywhere, and “Chewbacca Mom” was a guest on countless talkshows. In the meantime, more than 160 million users watched the video and were almost certainly unable to resist laughing along.



#DicksOutForHarambe

In May, the gorilla Harambe was shot dead after a four-year-old child fell into his zoo enclosure. Within hours, a video of the incident was shared several million times. Social networks were flooded with countless memes, and the hashtags #Justiceforharambe and #RIPHarambe began making the rounds. The comedian Brandon Wardell soon issued a call for “Dicks out for Harambe”. His appeal struck a chord with a growing number of people; some even put it into practice. It reached the point that Harambe was able to win 11,000 votes in the American election of November 2016.

Microsoft Chatbot

Microsoft has been working on artificial intelligence for some time. In order to gain insights into how people communicate with each other, they developed the chatbot Tay and set it loose on Twitter. It started out fairly harmlessly, but within a few hours Tay had turned into a racist, anti-Semitic, xenophobic misogynist. After 96,000 tweets, Microsoft pulled the plug. The question is, was the experiment a success?



Homewrecking penguin

National Geographic shared this heart-rending and dramatic film clip on Twitter. It shows a male penguin who returns to his nest to find another male at his mate's side. The protagonist attacks the interloper, but ultimately loses the bloody fight. The female dumps him, and the internet community weeps (or is simply disturbed by all the gore).



Social Media Party

Bento published an article on “This Spring’s 15 Most Unique Magazine Covers”. The German magazine Spiegel Online shared it on Facebook with the note “Cover 5 had us in tears!” and Vice commented, “So something unexpected happened at cover 2 and we started crying”. Something unexpected really did happen: A major meeting of all German social media editorial teams in the flurry of comments that ensued.



Photo: Laura DiMichele-Ross

Tom Hanks or Bill Murray

This photo fairly dated, but made a big comeback this year. Millions of internet users racked their brains to figure out whose visit had left this child so distinctly unimpressed. Was it Tom Hanks or Bill Murray?

Hydraulic Presses

There was no limit to this year’s orgy of destruction: slowly, steadily, and with frightening power, in countless videos hydraulic presses have crushed everything in their path, from bowling balls to a safe to something that had once been considered indestructible: the Nokia 3310.



Screenshot: youtube.com

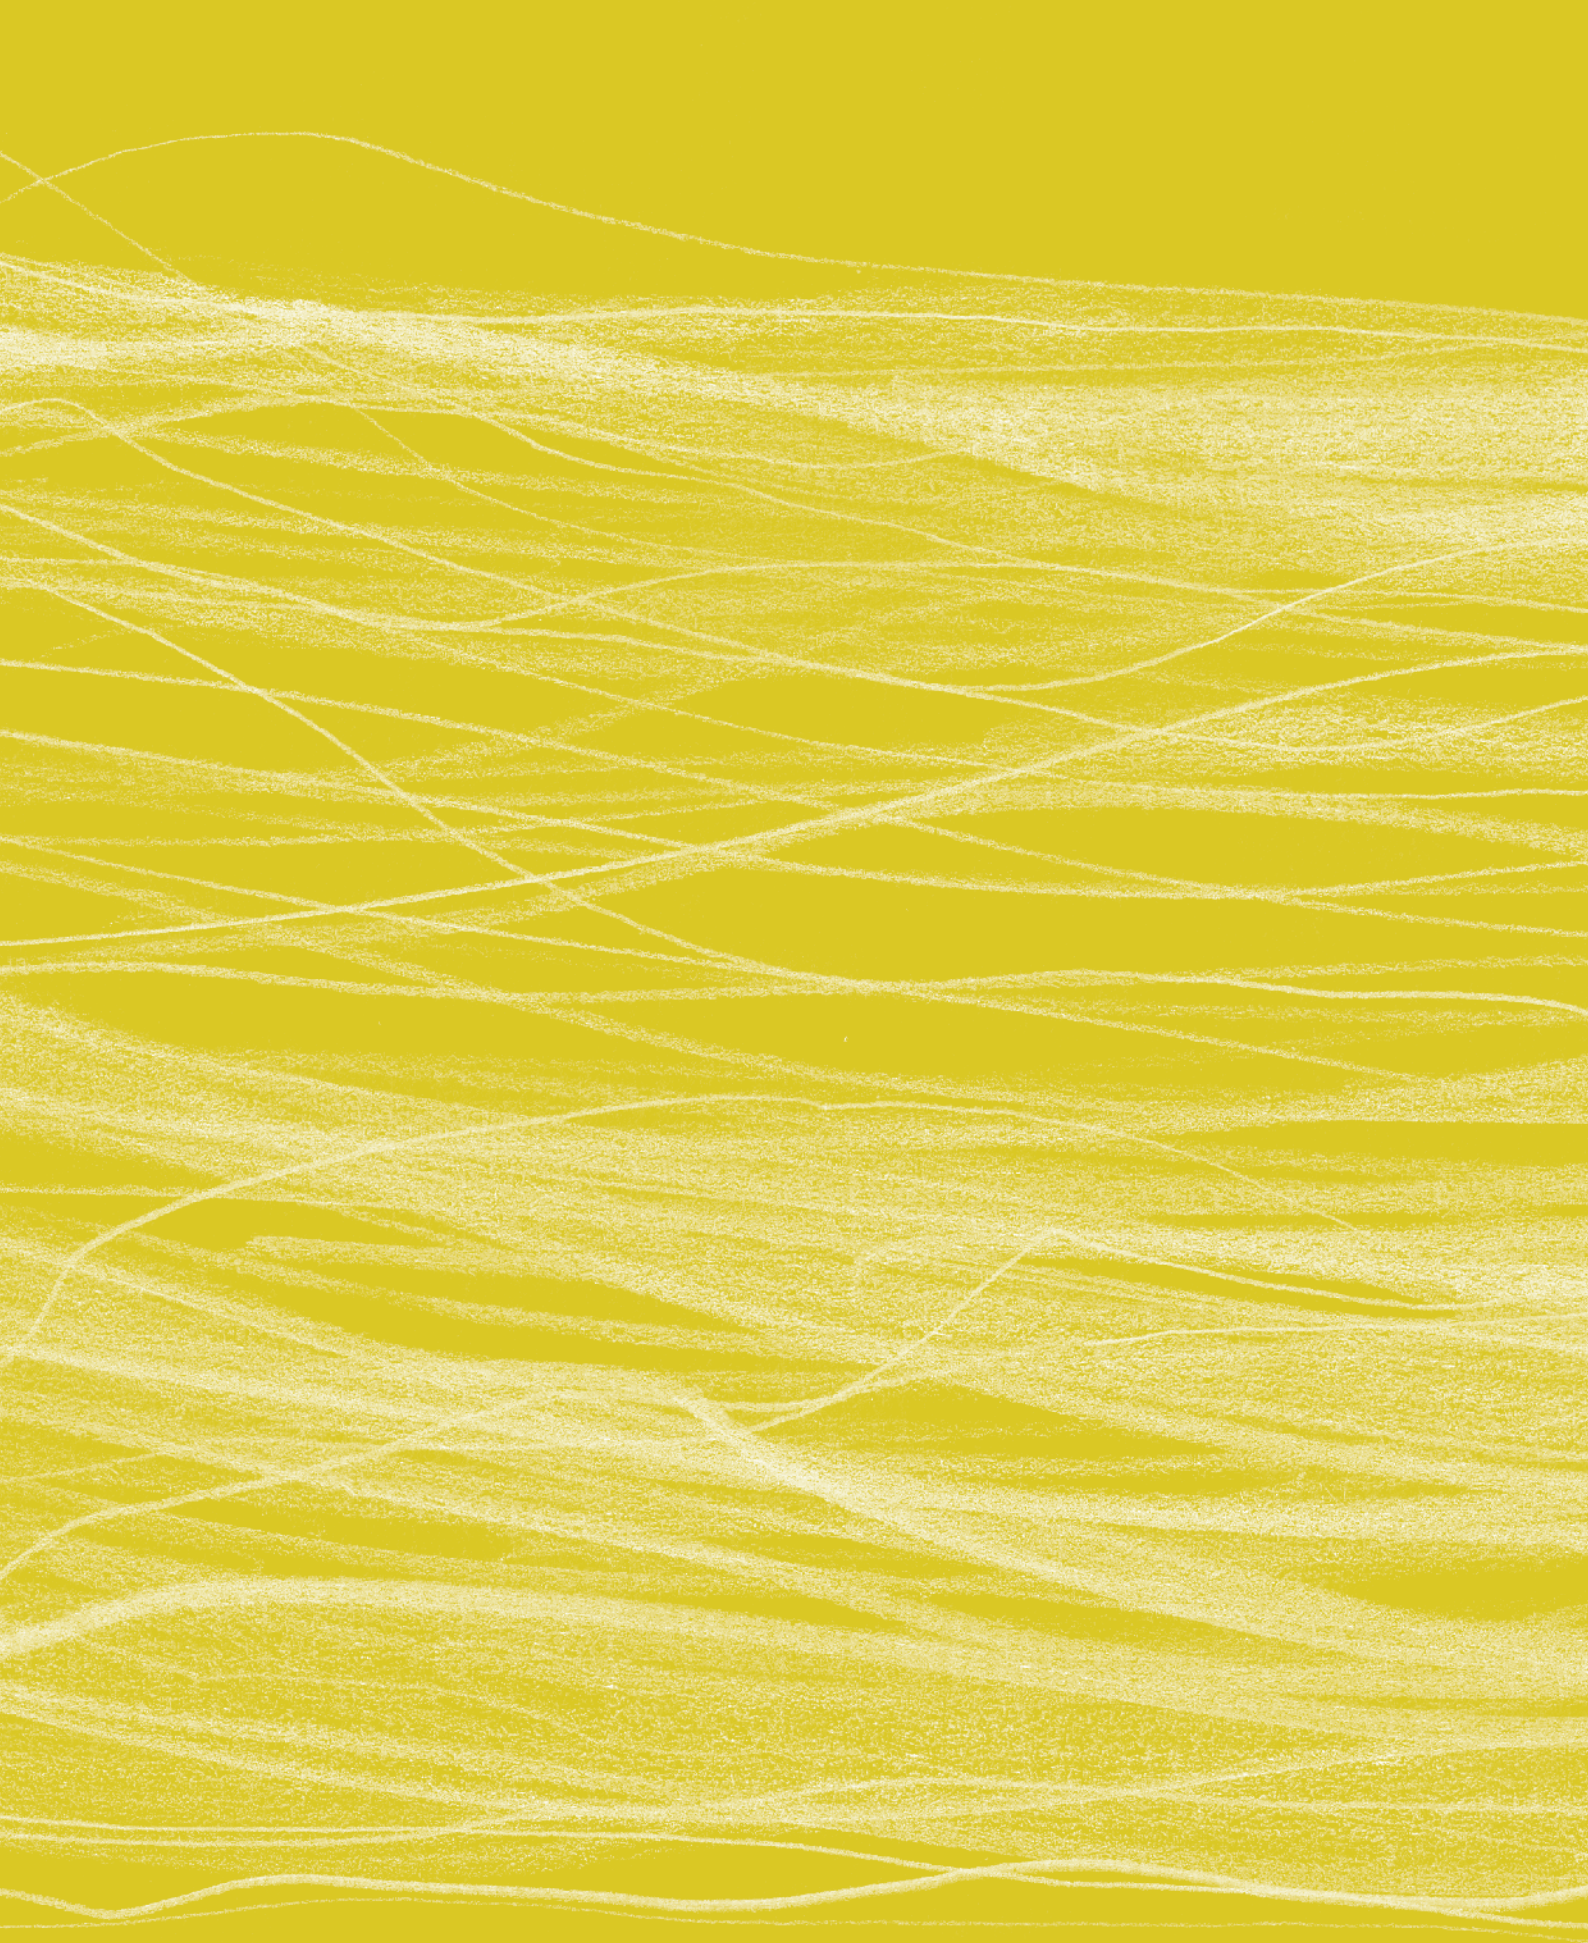


Photos (3): Karen Zack



Chihuahua or Muffin Puppy or Bagel Labradoodle or Fried Chicken ...

Chihuahua or muffin? Puppy or bagel? Labradoodle or chicken nuggets? This spring, the American Karen Zack asked us these questions and more. Some weren’t so easy to answer...



Life

Medicine

Internet of things

Artificial intelligence

Digital education

Catastrophes

Twitter trolls

Pokémon Go

Snapchat

Good sex

Serendipity



Predictive healthcare: Medicine in the data revolution

BY LYDIA HELLER



Apps and algorithms to help predict illness: Many of these applications fall into the “lifestyle and well-being” category of products, but they nonetheless indicate a trend which will change medicine. With big data, medical treatment will become more personalized, more preventative, more proactive.

Health apps for Smartphone are booming. Around 100,000 such apps already exist, meant to help with weight loss and to mitigate depression, to calculate fertility cycles, or to train the user in mindfulness. At the same time, new sensors are constantly being developed: fitness wristbands and smart watches count steps, monitor sleep and measure heart rates. Cameras, rings, patches and implantable sensors measure skin conductance, perspiration and blood values. Google, Apple, Microsoft, Samsung: in recent years all the big IT players have been bringing to market health applications for home use.

This is because lifestyle, fitness and health data applications have developed into a huge market in recent years. They form the missing piece of a puzzle that can perhaps make good on the promises of “personalized medicine” made a decade ago. At that time, the human genome had just been decoded. Using the genetic code, it was said that it would be finally possible to discover treatments for cardiovascular diseases, cancer, or Alzheimer’s. Success, however, has thus far remained elusive. Direct causal relations between genes and illnesses are hard to find, and our genome, so far as we know, works in a much more complex way than we had assumed.

Unbelievable volumes of data

Since then, not only has computer performance drastically improved and the cost of gene sequencing fallen enormously, but there are now unbelievable volumes of digital data available, gleaned from patient records, studies, and, not least, the plethora of health, lifestyle and fitness apps. As people collect more and more data on themselves, and as the number of connections and patterns emerging from this data increase, each individual can more precisely trace their own biological makeup.

“Just as the microscope made things visible which were much too small for the human eye”, wrote American economist Erik Brynjolfsson a few years ago, “the analysis of large volumes of data by means of algorithms makes connections visible which previously were far too big and complex for human understanding.” But lifestyle data, or the personal, health-related data collected by many fitness apps is not easy to relay and aggregate. At least, for the time being, not all of it is. Researchers worldwide are already working on programs that can reveal the complex relationships between body, environment and behaviour and simulate how patients will react to treatments, as well as assist in developing personalized medical interventions.

At the paediatric oncology clinic in Homburg, Norbert Graf is working together with mathematicians, molecular biologists and biological computer scientists to develop a computer model for Wilms’ tumour. This childhood kidney cancer, the professor explains, forces doctors to choose whether to operate immediately or to first treat the tumour with a course of chemotherapy in the hope of shrinking it, so as to render the surgery more straightforward. But not all children respond equally well to chemotherapy.

The program aims to generate a prognosis based on data about the

previous development of the tumour, medicines and their active ingredients and the widest possible range of clinical information on the patient. “We want to know how the tumour will respond to prior treatment.

US dollars for the *Precision Medicine Initiative*, which he inaugurated at the start of this year and which will see the genetic and health data of over one million Americans saved and made available for cross-referencing. This

On the one hand, the quality of data recorded by wearable devices and trackers frequently falls far short of medical standards. Studies have repeatedly shown that such devices can often generate false readings.

Ultimately the system should say: ‘the tumour won’t get any smaller, operate immediately’.” The bigger the volume of data on which the model can draw, and the more frequently its predictions can be measured against outcomes and adjusted accordingly, the more precise its prognoses will become.

Providing the best treatment right from the start

It would be immensely useful for doctors if it were easier to cross-reference data from medical records with personal information—and additionally with genetic test results and studies on the efficacy of different medications—, according to Norbert Graf. Many of his colleagues agree. “That way, we would be able to provide patients with the best treatment right from the start, and reduce the side effects they suffer.” Since 2011, clinics in several European countries have been working to network their databases, and to store information on, amongst other things, illness-related genetic and biological markers in blood and tissue samples. This has resulted in the the Biobanking and Biomolecular Resources Research Infrastructure (BBMRI).

In the USA President Barack Obama provided around 215 million

initiative should make it possible to perform tests in order to predict the effects of drugs. Analysis of this database should not only provide hints on how a treatment should be designed to battle an acute illness. The fact that this information is also linked to lifestyle data is “an incredible treasure trove” for medicine, says Norbert Graf, because it can also provide information on the likelihood of relapse.

Graf continues, “Following a successful course of cancer treatment, you always want to avoid a relapse. ‘Is there something special I should eat?’ is a common question, as is ‘Should I do more sport?’ And if I had, for example, information from this kind of health tracker about patients’ sports and nutrition, and if I had long-term information about who had or had not had a relapse—then I would be able to say to someone: ‘if you do this, or if you eat that, you’ll have such-and-such a chance of avoiding a relapse.’ We can’t do that yet.”

A data protection nightmare

Nonetheless, this development is a nightmare from the perspective of data protection. On the one hand, the quality of data recorded by wearable devices and trackers frequently falls far short of medical standards. Studies

have repeatedly shown that such devices can often generate false readings. On the other hand, critics fear that the storage of health data cannot be deemed sufficiently secure to guarantee anonymity. One fear is that this could lead to discrimination or disadvantages for those seeking employment, for example, should employers become aware of illnesses or predispositions to certain illnesses. Critics are also worried that in the future it could become obligatory for one to gather data on oneself using various trackers or apps, for the purpose of providing it to doctors or insurers.

Even now, insurers like the German public health insurance AOK or the Swiss Generali Versicherung have started rewarding customers with bonuses and discounts if they can prove they have a healthy lifestyle with data gathered by app. “Currently, it’s all voluntary”, says doctor and e-health expert Tobias Neisecke. “And it’s about rewarding someone who is being proactive about taking care of their health data. But it is probable that this could be turned around. At some point it will become about: ‘what’s my app score?’”

Health insurers insist that there is no disadvantage for members who decline to take part in this health monitoring. Nonetheless, though it remains an open question, bigger business will probably be made with the data itself; it will provide raw material for prognosis models which calculate health risks, not only with a view to creating treatments which are appropriate for target groups, but also for the purpose of developing preventative interventions.

Targeting and speaking early on with at-risk patients

Since 2014, the Carolinas HealthCare System, a network of doctors in the state of North Carolina, has looked at correlations between consumer data and health data in order to identify

patients who are at risk for specific illnesses. In Germany, the Elsevier Health Analytics think tank is working on algorithms which can look for patterns in anonymized health insurance data and identify groups of policy holders where there is a given probability that certain illnesses will arise. Doctors will be able to check their patient data against this filter and speak with at-risk patients early on.

The German health insurance provider AOK is also developing a “cardiovascular risk assessor”, according to Kai Kolpatzik from the AOK Federal Association in Berlin. It should predict “how high your risk is of having a stroke or heart attack over the next ten years, on the basis of age and blood pressure, whether you smoke, and your family’s medical history. And what’s exciting is that this can tell you things like: What will happen if I take this medication? What effect would a change in lifestyle have?”

Analysts calculate that if current double-digit annual growth figures persist, the market for personalized medicine will have a global turnover of 90 billion US dollars by 2023. This is money that should belong to the people who provide the data, says Ernst Hafen of ETH Zurich. Together with colleagues, he has initiated the MiData project: a co-operative whose members—patients and health professionals alike—are able to upload genetic and other health-related data onto a server, but decide for themselves what the data can be used for. Companies that use the data must pay for it. The proceeds are to be used to finance research projects which big private firms see as unprofitable.


Apart from the question of who will carry out medical research in the future and who will benefit from it, the predictive analysis of this data is bound to change medicine: instead of diagnosing acute illness, the question is increasingly one of predicting the likelihood of problems occurring down the road. “We are no longer just sick or healthy”, says the medical

ethics expert Peter Dabrock, “we are the carriers of given risk profiles. And that’s where it becomes ethically and economically interesting, because that poses a whole new array of questions in terms of the consequences that this has for health insurers. Today, we say: carriers of a given genetic mutation, for example, have a claim for a given treatment, which we pay for. Soon, it could be: We’ll pay for a treatment with 70 percent chance of success. But what about 65 percent? Will we still pay for that?” ■



Photo: private

Lydia Heller is a freelance writer, reporter and presenter, mainly working with Deutschlandradio Kultur, Deutschlandfunk and Deutsche Welle. Since 2008, her favourite—but not her only—job has been writing radio features about the environment, technology and science.



Attack of the fridges

BY JESSICA BINSCH

The networking of everyday objects is speeding ahead. From toothbrushes to baby monitors, all kinds of gadgets are getting connected to the internet. But the internet of things can be hacked, and botnets made of toasters can take over our machines.

When looking to buy a new home appliance, you normally wouldn't give much thought to hacker attacks. But the next time you're shopping, maybe you should keep Andrew McGill's toaster in mind. McGill is a programmer and journalist; he works for the American magazine *The Atlantic* and his toaster was recently hacked.

Luckily, it wasn't McGill's actual toaster. But it should still give us cause for concern. McGill had simulated a toaster for an experiment—a toaster with an internet connection. He wanted to find out how quickly the gadget would be targeted by hackers. McGill was “fully expecting to wait days—or weeks—to see a hack attempt”, as he wrote in his report for

The Atlantic. In fact it took less than an hour. Within the first twelve hours there were a further 300 hacking attempts.

McGill's experiment is more than just an amusing anecdote. More and more everyday items are connected to the internet. From baby monitors to toothbrushes—all manner of gadgets are becoming “smart”. Experts predict that the market for networked gadgets will soon be worth billions of dollars annually. No wonder, then, that more and more companies are looking for a piece of the action. Internet giants Google and Amazon have brought their own control centres for networked households onto the market. Google Home and Amazon Echo react to spoken instructions from their users

via microphones and built-in software assistants.

Even small and medium enterprises assume that in a few years practically all household goods will at least have the option of going online. We can observe the same development with television: there are now hardly any television sets for sale which are not smart.

But in the scramble for the market, security is falling by the wayside. It is becoming more and more clear that networked devices have their vulnerabilities, and 2016 could be a turning point. This past year, the first massive internet attack associated with networked gadgets was made public.

One Friday in October, internet users in the USA faced massive

network failures. Big online services like Netflix and Spotify went down, as did sites like Reddit, the New York Times or Wired.

Among the culprits were insecure webcams. Hackers had joined millions of devices together into a botnet. This botnet targeted the DNS provider Dyn. Companies like Dyn are responsible for translating website names into IP addresses, the only way that a browser can call up the required site. Dyn is the internet's telephone directory—and a weak spot in the global infrastructure.

The company was overwhelmed by a massive wave of nonsense requests, in other words, a classic DDoS attack, which bring servers to their knees by overloading them. For attacks like these, attackers use botnets made up of devices which they have brought under their control. Until now, this generally only meant computers and laptops, not video recorders and webcams.

Experts had already been warning for some time that networked devices could be used for attacks. The IT journalist Brian Krebs experienced this first-hand, when his website was attacked by a botnet made up of surveillance cameras and digital video recorders. The software employed was amateurishly simple, but its effect was devastating.

Warnings are growing louder. “We need to save the internet from the internet of things”, declared IT security expert Bruce Schneier in the technology magazine Motherboard. Schneier issued his call to arms only a few weeks before the massive attacks at the end of October. In hindsight it was almost prophetic.

The problem lies within the networked devices themselves. Or rather, with their manufacturers. Companies construct their products often without any thought of security and maintenance, says Michelle Thorne. Thorne works for the Mozilla Foundation, which is behind the Firefox internet browser. She has written a book together with Peter Bihr about the internet of things, called “Understanding the Connected Home”.

“People buy a fridge, and then at some point they have to update it”, says Thorne. “But the tech companies are not ready to

support that or think about long-term maintenance.”

Often, updates are not possible, nor there are provisions for changing the standard password. This was how the attack on Dyn in October 2016 took place: the hackers used surveillance cameras from a Chinese manufacturer, which were running with a known standard password. Not all companies are familiar enough with internet security to properly secure the networked devices they started building. No one knows exactly how many cheap surveillance cameras or video recorders are connected to the internet without proper safeguarding.

There is hope that the recent attacks on the infrastructure of the internet will at least have one positive effect. The problems are now known, the wide-ranging impacts of security flaws have been comprehensively demonstrated. That has brought state regulators onto the scene. The German authority for IT security, the Federal Office for Information Security (BSI) is now calling on manufacturers to do better.

The majority of household goods connected to the internet are “insufficiently protected against cyber attacks when they arrive from the factory and can therefore be easily taken over by attackers and put to criminal use”, warns the BSI. “We therefore require that manufacturers of networked goods improve the security of their products and that, when developing new products, they look not only at the functional and price aspects of the item but also at the necessary security aspects.” Manufacturers should encrypt internet communication and provide updates.

Experts are also discussing ideas for an IT quality seal. Such labelling would inform consumers that products meet certain safety standards. Whether stronger rules are required is still up for debate. And even if they are, it could take some time before they are in place.

It could indeed be that security becomes a sales angle for networked

devices. That may be an optimistic scenario, but it is not inconceivable. A similar development led to a change in messenger apps. Only a few years ago, security in chat services was a niche topic, addressed only by a few small providers. Then the giant Whatsapp began encrypting its users' messages. A major impulse behind this were Edward Snowden's revelations of widespread of digital communications surveillance.


It is possible that the massive DDoS attack of October 2016 will make people more careful when buying. Manufacturers will be placed under greater pressure to make their networked products more secure. In any case, the market is very diverse: not all companies offering networked devices are necessarily versed in IT security. It is likely that the incident in October was not the last time internet-enabled household goods will play a part in a cyber attack. ■

2016 could be a turning point. This past year, the first massive internet attack associated with networked gadgets was made public.




Photo: private

Jessica Binsch works as a freelance journalist in Berlin and reports on digitalization and society. She is especially interested in internet politics, internet activism and the social impacts of technological developments.



Artificial intelligence: The dreaming algorithm

BY CHRISTOPH DROSSER



Sometimes knowledge hides away in difficult places, but now and then the time is ripe to venture out in search of it, no matter how hard the journey. Welcome to an expedition, an ascent, into the rarefied world of machine learning.

Base Camp

Don't set off without packing the following basics

A computer, it is often said, only knows as much as the programmer that has given it its instructions: all it does is follow instructions. This is true of the simplest levels of machinery: software works on the commands from the programmer, going through them line by line. But does that mean that a computer can't learn? To say that would be just as false as to say that a pupil can never be smarter than their teacher. So, just as a good teacher doesn't just let his pupils learn facts by rote, but nurtures their own development, a computer can be programmed so that, the more time it devotes to fulfilling its tasks, it continuously improves in its ability to do so. Welcome to the world of machine learning.

The first self-teaching program to make a splash was developed by the IBM researcher Arthur Samuel in 1956. The software played draughts at a respectable amateur level. At the start, the computer only knew the rules of the game and a few rules of thumb that Samuel had given it. But with every game, the machine learned more. After eight to ten hours of training time, it was better than its creator. Today, humans can no longer beat computers at draughts. In chess, the computer is at least an equal match for us, and since Google's AlphaGo program beat the European Go champion, humans are no longer undefeated in any board game.

Machine learning is a subdomain of Artificial Intelligence. Today, a wide

variety of software techniques fall under this category: computers learn how to identify humans in photos. They drive driverless cars through city traffic, after they have trained for a few thousand hours. They find patterns in big data.

In many of these learning techniques, a human is still the teacher: the human sets a goal and evaluates the computer's performance, while the computer varies and adjusts its behaviour in order to get better marks. At the same time, what could be called unsupervised learning plays an important role: the computer has to make sense on its own of masses of data. Thus, Google feeds millions of photos into a computer network, and the program creates automatic categories like "cat" or "human". This closely resembles the way a young child learns, as they create categories before they can name them.

First climb

Let's go! On the gentler slopes you will encounter knowledge which can bring you out in a sweat.

As a first climbing exercise, let's play a game which is already too simple for five-year-olds: Noughts and Crosses. The board is made up of squares arranged three by three. Two players take it in turns to set down their counters. Whoever manages to get three counters in a row, straight or diagonal, wins. There are 255,169 possible outcomes in this game. In 131,185 of them, the player who goes first wins, with the second player winning in 77,904

variants. 46,080 variants end as a draw. More important than this is the fact that a “smart” player will never lose a game: Regardless of whether they go first or second, they can set down their pieces (or draw their noughts or crosses if playing with pen and paper) so that the game at least comes out as a draw.

How can you figure out the best move to make in a given situation? In Noughts and Crosses, all possible moves can be calculated beforehand. That leads to a decision tree: a player looks at all the moves that they can make given the current state of play, then at all possible responding moves from their opponent, and so on. In chess, this leads to an explosion in the number of possible configurations; but in Noughts and Crosses, the potential combinations are limited enough to be manageable: after at least nine moves, the playing field is full and will show any of 138 end positions. Every branch ends with the victory of one of the players, or a draw.

In order to assign a value to every playing position, one evaluates every leaf on this tree: a win gets a value of +1, a loss gets -1 and a draw is given as 0. Then take a step back through the

And from a completely ignorant program, we get one that never loses a game.

game. Every sub-branch of the decision tree is allocated a playing position and a value, which is the highest of the following values if it is your turn, and the lowest of the following values if it is the other player's turn. At the end, all positions have an evaluation of 1, 0 or -1. Branches with a value of 1 mark a strategy which can only win.

Breathe deeply

An example: let's assume that our opponent plays first and places their cross in the middle of the board (the best starting move). We then place our nought either in a corner square or in a square in the middle of one of the grid's sides. Which of these moves is better? Let's look at variants in which we choose the middle of the left-hand row. There are then four essentially different possible responses for the opposing player to choose from. Let's assume that they place their cross directly above our nought. Then in our next move we have no choice: We must place a nought in the lower right-hand square, in order to stop a diagonal line from being created. Then, the opposing player can knock us out of the game with a cross in the middle of the upper row.

In fact, our first move was fatal. It leads to a -1 in the decision tree, and should be avoided. If we had put our nought in the corner on our second move, even against the smartest player we would have an opportunity to fight them to a draw. This move has a value of 0.

How could we get a computer program to play with this strategy? First possibility: all the values in the decision tree are put in a table. The computer looks at every move in its table and chooses the move with the highest value. It plays perfectly from the first move and has no need to “think” at any point. Second possibility: The computer starts the game totally “stupid”. In every situation it marks the possible moves with the values -1, 0 and 1. As soon as a game is over, it changes these values retrospectively, in light of the outcome. In this way, its evaluation of the game-play will constantly improve.

If we now let the computer play against itself, something interesting happens: While both parties (which are in fact just one party) have no idea about the game, their retentive memory helps them to try out the different possible moves and learn which

approach is good for one position and which is bad. And from a completely ignorant program, we get one that never loses a game.

Steep slope

Take deep breaths! It's not how you expected—but you'll make it.

Board games are comprehensible worlds with clear rules and unambiguous situations. While people can quickly surrender in the face of their complexity, for computers they are straightforward. On the other hand, thinking through muddy reality, which is easy for us humans, is extremely difficult for computers. Take, for example, an exercise which most people would hardly even label thinking: classification. Is that a photo of a cat or dog? Is that the voice of mother, or a stranger? Is that thing in the road a plastic bag or a rock? We are able to arrive at the right answers without any real thought and with an astounding degree of accuracy. But even we don't know exactly how we manage it.

In the 1970s and 80s, people tried to teach computers to classify things using rules developed by experts: a cat is an animal with pointed ears and whiskers; a mouse is grey and has a long tail. This method didn't at all work well. In recent years, we have had much more success with what is called neuronal nets, which imitate the structure of the human brain. They perform astoundingly well with large volumes of data.

Neuronal nets were actually invented in the Fifties, but they only came into their own with the development of modern computing power, under the label “deep learning”. William Jones and Josiah Hoskins described a very simple example in 1987 in *Byte* magazine. The neuronal net should help Little Red Riding Hood to survive the deep dark wood. In particular, it should keep her from

being eaten by the wolf. The story also features grandma, and a huntsman, who saves Little Red Riding Hood.

Big ears, big eyes, big teeth

The program doesn't know humans. It only sees particular physical characteristics and has to derive a particular approach from them. The wolf has big ears, big eyes and big teeth. When Little Red Riding Hood meets him, she should run away, scream, and look for the huntsman. Grandma has big eyes, wrinkles and is friendly. If Little Red Riding Hood spies her, she should come close, kiss her on her cheek, and offer her the food she has brought. The huntsman has big ears and is friendly and attractive. The desired behaviour: Little Red Riding Hood should approach him, offer him food and flirt with him (the article is, as I've said, almost 30 years old).

We can see right away that the relationship between sensory impressions and desired behaviour is far from straightforward: A being with big ears could be the wolf, but also could be the huntsman, and these each require a very different reaction.

The neuronal net is made up of two "layers" of cells: It has six input cells, which note the major characteristics of the actors (big ears, big eyes, etc.) and seven output cells, which correspond to Little Red Riding Hood's repertoire of potential behaviours (running away, screaming, looking for the huntsman, etc.).

Every input cell is linked to every output cell, and at the start, each of these connections has a given "weight"—a number that describes its strength. We start with relatively small, randomly-chosen weights. This initiates the self-training of the network. It is fed successively with the input values for wolf, Grandma and huntsman (the first figure stands for "big ears", the last for "attractive"):

Wolf: (1, 1, 1, 0, 0, 0)
Grandmother: (0, 1, 0, 1, 1, 0)
Huntsman: (1, 0, 0, 1, 0, 1)

Reaction to the huntsman:
(0, 0, 0, 0, 1, 1, 1)

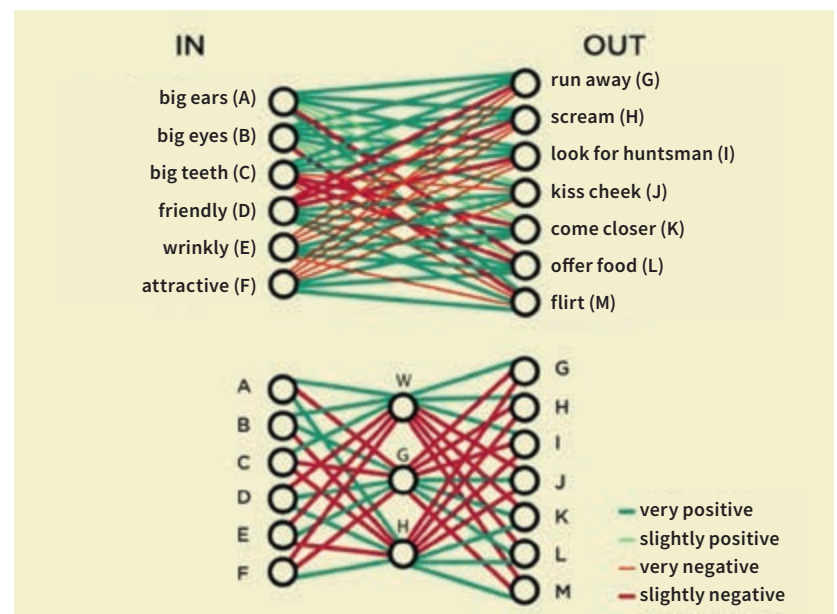
The corresponding input value is passed from the input cells to all output cells (from "run away" to "flirt"), and is this multiplied by the weight of the respective connection. For each of the seven task neurons (from "run away" to "flirt"), six numerical values are given, which are added together. If the sum exceeds a threshold (e.g. 2.5) then the neuron "fires"—and the output cell assumes the value 1.

At the start, the net behaves randomly, because the weights of the connections are chosen at random. So that it can learn, we must compare the result with the desired action from Little Red Riding Hood:

Reaction to the wolf:
(1, 1, 1, 0, 0, 0, 0)
Reaction to the Grandmother
(0, 0, 0, 1, 1, 1, 0)

and alter the strength of the connection on that basis. After about 15 run-throughs, the net becomes largely stable. It develops the connections shown below left.

Why create this complicated training program, though, when we already know all the rules? In practice, the net is used in situations where the desired output is only known for a limited number of training examples. If the net is to analyze photos of animals (as digital volumes of pixels), and learn from them how to name the animals, we don't say that a cat has pointed ears. That would mean that when the net has correctly identified the animal, it would not be able to formulate why it described a given image as a "cat". Rather, it can re-use what it has learned on new pictures and recognize cats there too.



This example shows how a neuronal net learns. The graphic above shows the net after 15 training steps. The connections between IN and OUT have assumed positive or negative weight, so that Little Red Riding Hood can react correctly to the other party's characteristics. In the simulation shown below, three additional neurons are added. They specialize in the recognition of the wolf (W), Grandma (G) and huntsman (H).

The neural net has no prior knowledge and extremely limited tact. Software engineers need to train their algorithms in greater sensitivity.



Photo: Liesa Johansson

Christoph Drösner also wrote this article for his new book "Total berechenbar: Wie Algorithmen für uns entscheiden" (Totally calculable: How algorithms are making decisions for us), using a computer. Despite all progress in Deep Learning he remains sceptical that a computer will ever be able to write such things by itself.

A strong drive to flirt

We have trained the Little Red Riding Hood net on three examples. There are a total of 64 possible inputs for the network, from (0, 0, 0, 0, 0, 0) to (1, 1, 1, 1, 1, 1). And each of these inputs will create an output in the net. Is this plausible?

For example, we can imagine what would happen if the wolf put on sunglasses and started being really friendly. That would correspond to the input values (1, 0, 1, 1, 0, 0). The output of the net which has been trained here would be: a certain tendency towards the correct reaction to the wolf (running away, screaming, looking for the huntsman), but also a strong drive to flirt. Clearly the wolf presenting himself like this confuses the girl, which is also understandable. Ambivalent input creates ambivalent behaviour.

Onto the summit

Now it's getting drafty: You must master this theory if you want to rise to the occasion.

In order to further increase the performance of neural nets, the developers have come up with a trick: they insert a "hidden" layer of neurons between the input and output cells. Where the net is correctly trained, these neurons develop certain specializations. In our example, three cells can be inserted in the hope that these will specialize in the recognition of the Grandmother, the wolf and the huntsman (W, G and H in the graphic on the right). In the experiment they operate without any help. Cell W reacts especially to inputs which correspond to characteristics of the wolf, and triggers an appropriate response. The invention in 1986 of this hidden layer and its reasoning processes (so-called back propagation) marked a breakthrough.

This layer can be seen as an ever-higher level of abstraction of the

sensory input: a net which has to recognize images only looks at disaggregated parts of images at the input level. The first hidden level of neurons will, perhaps, recognize starkly-contrasting edges. That is the basis for identifying, for example, circles or squares at the next level. Deeper in the net, neurons develop which can, for example, recognize eyes or even a cat's head.

Sometimes the net also gives results which its creators rightly find embarrassing. For example, an automated image recognition program used by the photo service Flickr categorized men with black skin as "apes". The gate of the Dachau concentration camp was labelled a "climbing frame". The neural net has no prior knowledge and extremely limited tact. Software engineers need to train their algorithms in greater sensitivity.

Deep Learning is now yielding successes which eluded artificial intelligence for decades: the nets can recognize human faces on photos with confidence. They can understand spoken language very well. Skype can interpret between speakers of different languages in real time.

For the learning programmes named here, there was always a human teacher which trained the program in the correct answers. But increasingly, these nets are learning independently. They are fed huge volumes of data, and left to make sense of it themselves. Google engineers caused a stir two years ago when they put neural net "on drugs". If you require the net to find an object in a plain image, as when a person looks for patterns in clouds, it will hallucinate and see, for example, fantastical fishes in the sky where there are none. The machines have learned to dream. ■

The article first appeared in ZEIT Wissen No. 5/2016, 16 August 2016. Reprinted with friendly permission from the Zeit Verlag.

Premium Interconnection Services. Worldwide.

DE-CIX has been offering Premium Interconnection Services for more than 20 years. More than 1000 networks from over 60 countries connect at our Internet Exchanges in Europe, North America and the Middle East.

Join us now and connect to our Internet Exchanges with hundreds of international and regional networks to enhance the quality of your IP traffic, reduce latency, gain control over your routing and massively improve the end-user experience.

**Where
networks
meet**



What happened online?

January 2016

4

04/01 The CDU politician Andrea Voßhoff is now Germany's Federal Commissioner for Data Protection and Freedom of Information (BfDI), an independent supreme Federal authority. While enjoying freedom from legal or administrative supervision at the hands of the federal government, the Commissioner is still not able to issue sanctions.

6

06/01 The German Federation of Consumer Organizations (vzbv) rebukes Google for its new data protection declaration. Google reserves the right to analyze users' emails, amongst other things, in order to personalize the advertising they see.

7

07/01 The Berlin District Court rules that a Facebook account can be bequeathed in a will. Facebook is obliged to give the parents of a dead girl access to her account.

11

11/01 The German Federal Intelligence Service (BND) resumes cooperation with the American National Security Agency at the Bad Aibling field station. The American secret service will continue to hand over its search terms (selectors), but must now be able to justify them. Until now, no request from the NSA had been declined.

13

13/01 According to a ruling by the European Court of Human Rights (ECHR), the Hungarian Surveillance Act is in breach of the Convention on Human Rights. The law allows, among other things, for every person's communications to be individually monitored in Hungary.

14

14/01 Ernst Uhrlau, former head of the BND, speaking before the German parliamentary committee investigating NSA surveillance practices, expresses doubt regarding the statement of the former Chancellery Chief of Staff, Thomas de Maizière (CDU), that he was not informed about the termination of the BND-NSA collaboration "Eikonal".